EVLA Memo #181
Atmospheric Phase Interferometer Instrument Description

Keith Morris

September 16, 2014

1 Introduction and Background

The Atmospheric Phase Interferometer (API) is a two-element atmospheric seeing monitor located at the Very Large Array (VLA) site. The instrument measures turbulent refractive index variation through the atmosphere by examining phase differences in a satellite beacon signal detected at two (or more) antennas. Instability of the atmosphere above the VLA causes phase fluctuations in propagating radio signals, which scale with frequency.

The VLA is dynamically scheduled, in other words the choice of what to observe next depends on the observing conditions: wind speed, atmospheric phase stability, and other priorities. In order to know which frequency to observe at any given time, the atmospheric phase stability at the VLA is measured continuously. If fluctuations are above predetermined thresholds, then observing at the higher frequencies will not be undertaken at that time.

From 1997 until 2014 this continuous measurement was made with the Atmospheric Phase Interferometer (API): a two-element atmospheric phase monitor located at the VLA site. The instrument measures atmospheric phase by examining phase differences in a satellite beacon signal at ~12 GHz measured at four antennas separated by 300 m and recording those phase differences as a time series. The phase differences give the continuous measurement of atmospheric phase stability above the VLA that is required for dynamic scheduling.

The first dedicated API to be built was deployed at the NMA. This API used a geosynchronous satellite signal at 19.45 GHz for its measurement, on a baseline of 35 m. Frequencies of 10-20 GHz and baselines of 10-300 m have been used for all subsequent APIs. The choice of frequency for a particular API is driven mostly by (1) availability of satellite signals, which are common in the 10-20 GHz range, especially around 12 GHz, (2) a desire to not go too high in frequency, which avoids some difficult instrumentation, and (3) a desire to not go too low in frequency, to be sure that the measurement is still dominated by tropospheric, and not ionospheric, fluctuations.

The API at the NMA was followed by an instrument deployed at Mauna Kea, Hawaii, to measure the atmospheric stability in preparation for the new SubMillimeter Array (SMA). Shortly thereafter, NRAO built and deployed an API at the future ALMA site, and another at the Very Long Baseline Array (VLBA) antenna site on Mauna Kea, to measure and compare atmospheric stability at those two locations (Holdaway et al. [1995]). The design and construction of the NRAO APIs is described in detail in Radford et al. [1996]. Another API was deployed at an alternative possible ALMA site near where that telescope was eventually built.

The instrument at the VLBA site on Mauna Kea was deployed to the VLA site in 1997, as a full-time atmospheric stability measuring instrument (Carilli et al. [1998]). Following the practice of the previous APIs, the VLA API measures a continuous time series of phase differences on a 300 meter baseline once per second, and processes each ten minutes of this time series in real time to produce a continuous measure of atmospheric phase stability above the VLA. Ten minutes is

considered sufficient because enough independent parcels of air will have blown over the two antennas of the API to obtain an adequate statistical measurement of the atmospheric stability, even under very calm conditions. The VLA API has been operating continuously, with various modifications, since its original installation. While it has performed admirably, it is beginning to suffer from reliability problems.

In March 2011 we undertook the development of an upgraded instrument design, one that would be reliable enough to produce data for the expected lifetime of the VLA telescope (~2030), and address some of the limitations of the original system. This upgrade project was also an opportunity to design features into the instrument that will allow us to study the directional nature of the turbulence, probe the distance structure function, and to study the instrument itself.

The API measures the phase of an extremely weak (picowatts as measured at the ground) satellite signal that has propagated through the atmosphere and been detected at two (or more) locations on the ground. In order to do this, the instrument must:

1. Detect the weak microwave signals
2. Translate the detected signals to a frequency convenient for transport over distances up to 1000 meters
3. Filter unwanted signal components
4. Translate the signals to a frequency convenient for digitizing
5. Digitize and analyze the signals.

Based on funding, and a requirement to prove the design (see below), the API upgrade project was split into two "phases." The first ("Phase I") would include design and construction of most of the central electronics, but would include only two elements of the full 4-element API. These two elements would be tested while co-located, within 5 meters of the existing API. The second ("Phase II") would comprise the completion of the full 4-element API.

2 Technical Requirements for the Upgraded Instrument

In addition to the general requirements listed above, the primary capabilities which the new instrument must have are

1. Flexibility in frequency coverage – when the original satellite with the 11.3 GHz beacon was decommissioned about 6 years ago, the API had to be modified to allow frequency conversion from a different beacon frequency. The modifications necessary to allow this in the existing API are mostly improvised from available parts, and some portions of this scheme are aging and prone to failure. The new API shall have remotely tunable oscillators, which will allow for adaptation to other beacons.
2. Redundancy and diversity of measurement length scales and directions – the baseline directions of the VLA antennas span the horizontal plane, and as such could be sensitive to path length variation that is not detected by the current one-dimensional API measurement. Like the VLA, the new API shall span the plane, and even if we find no particular direction dependence to the turbulence measurement, the extra baselines provide redundancy in the event of a detector failure. Redundancy increases the reliability of the instrument.

3. A factor of two or more increase in sensitivity over the old instrument. That is, the new instrument shall be able to resolve changes in saturation phase of 0.4 degrees or less. This requirement arises from the VLA frequency-based observing condition thresholds. The highest-frequency receiver band has a threshold of 5 degrees root-mean-square (RMS) or less, while the next-lower receiver band has a threshold of 7 degrees RMS or less. The new instrument must reliably detect the difference between these two thresholds. For an unambiguous detection, we require that difference to be 5 standard deviations or more; therefore the instrument must have a maximum measurement uncertainty of 0.4 degrees RMS. Note that the existing API has a measurement uncertainty of around one degree RMS in the saturation phase.

The locations of the receiving antennas relative to the electronics shack drove an additional constraint: the signal attenuation due to the 900 m cable haul from the most distant antenna to the shack made the use of coaxial cable for transmitting the intermediate frequency signal undesirably lossy. Instead, fiber optic cable provides a lower-attenuation signal transmission, where signal attenuation is dominated by the modulator/demodulator and not by the link length.

These top-level technical requirements drive requirements for the subsystems, namely the phase sensitivity of the instrument, the stability of the oscillators, and the choice of local oscillator and intermediate RF frequencies.

2.1 Sensitivity

Although the smallest saturation phase threshold that is used in the dynamic scheduling algorithm is 5º for the 40-50GHz (Q) observing band, we must be able to resolve the difference between 5º and 7º, which is the next-lowest threshold, for the 26-40GHz (Ka) band. An instrumental measurement sensitivity of 1º RMS allows only a 2-standard-deviation detection of the difference between these two thresholds. A measurement sensitivity of 0.4º RMS would yield a 5-standard-deviation detection.

The satellite signal itself, as detected on the ground, has a Signal-to-Noise Ratio (SNR) of approximately 23dB; in other words, a factor of about 200 separates the amplitude of the sinusoidal beacon signal from the inherent noise floor. This means that the zero-crossing point of the sine wave has an inherent 0.5% uncertainty, and this uncertainty propagates to the phase measurement. Each sinusoid's noise can be treated as independent. If this were a simple zero-crossing phase detector, such as an oscilloscope, the theoretical uncertainty e_{total} in the resolution of the phase angle of two such signals, e_1 and e_2, would be

$$e_{total} = \sqrt{e_1^2 + e_2^2} \times 360^o = \sqrt{2 \times (0.005)^2} \times 360^o = 2.5^o \qquad (19)$$

which fails to meet even the coarsest resolution requirement. Fortunately, we can do much better than that. The API uses an averaging complex correlation to measure the phase. The value of the zero-crossing location is the mean across $1/f_{baseband}$, or about 10,000 samples. The uncertainty in the zero-crossing location is reduced by a factor of $\sqrt{10000} \sim 100$. This means that the SNR of the signal does not limit the phase resolution as it would in a direct zero-crossing detecting system.

Given the LNB manufacturer's stated noise figure, and the noise and gain figures of all subsequent devices, we can predict the total noise figure of the system. Table 1 shows the gain and noise contribution of each amplifier in the system, and the gain and noise budget for the

system as a whole, calculated from manufacturers' device data. Clearly the gain of the LNB is high enough and its noise figure low enough, that the system SNR is fixed when the signal leaves the LNB. Additionally, amplifiers have been chosen such that they remain in linear operation, free of gain compression, over the range of expected signal levels.

Table 1: System gain and noise budget, calculated from manufacturers' device data. LNB and intermediate amplifiers are shown.

Device	**Location**	**Gain (dB)**	**Noise Figure (dB)**
LNB	Front end	60	0.6
ZFL-1000LN	Front end	25	2.9
Fiber link	1st conversion	-35	3.0
2nd conversion mixer	2nd conversion	-7	7
ZX60-2531M	2nd conversion	33	3.5
ZFL-1000LN	2nd conversion	23	2.9
ERA-5SM	3rd conversion	12	4.3
3rd conversion mixer	3rd conversion	-7	7
Total		105	0.6

2.2 Frequency Conversion and Filtering

The microwave frequency of the satellite signal is a convenient, industry-standard frequency for communications and television distribution – it has adequate bandwidth for the signal payload, the electronic components are compact and lightweight, the transmission through the atmosphere is well-understood. But this frequency is less than convenient for other types of signal processing, such as narrow bandwidth filtering, modulation of lasers, or precision phase-detection. The API detects the phase of an extremely narrow bandwidth sinusoid, and the high frequencies that allow for wide-bandwidth signal transmission add unnecessarily to the cost of the components. For these reasons we convert the frequency of the signal, as early in the signal chain as possible.

The process preserves the phase of the original signal; any subsequent frequency conversion operations will also preserve the phase of the original signal, as long as the phase of the LO signal is controlled. The result is a signal, at some arbitrary frequency, with the desired phase.

At each stage of the three stages of frequency conversion, we perform band-limiting filtering. This reduces pressure on the gain cascade – the amplifiers are more likely to remain in their linear regions if they are not required to amplify unnecessary signal bandwidth. It also removes unwanted signal components that threaten to corrupt the signal multiplicatively. **Error! Reference source not found.** depicts the spectra of the signals and the filter for a typical API frequency conversion. In each of the three frequency conversion in the API, a filter selects the difference frequency, while rejecting the LO, RF, and sum frequency.

Table 2 lists the three frequency conversions – the RF, LO, and IF frequencies, and the bandwidths and types of the filters used at each stage.

Table 2: Conversion frequencies and signal component selection filters

Conversion	LO	RF	IF = RF - LO	Filter frequency/width/type
1st	10.75GHz	11.701GHz	951MHz	950MHz/12MHz/ bandpass
2nd	961.01MHz	951MHz	10.01MHz	10MHz/1MHz/ bandpass
3rd	10MHz	10.01MHz	10kHz	10kHz/20kHz/ lowpass

2.3 Oscillator Frequency Stability

All local oscillator signals in the API are phase-locked to the 10 MHz LO reference, or else employ the 10 MHz oscillator signal directly, as in the case of the 3rd frequency conversion. The 10.75 GHz LNB oscillator and the 961 MHz 2nd LO are each coherent with each other and among the four antennas.

The two primary source of jitter in the system are the 10 MHz reference oscillator, because its jitter is multiplied by a factor of 1075, and the LNB 10.75 GHz oscillator, because it is the highest frequency signal generated in the system. The LNB 10.75GHz oscillator contributes approximately 0.1fs (10^{-16} seconds) of jitter. Since it is phase-locked to the reference signal, any jitter present in the reference signal that is within the phase-locked loop bandwidth will also appear in the output signal. In this case, the 10MHz reference oscillator contributes 3 orders of magnitude less than that, but is multiplied in frequency by a factor of 1075, so appears on the same order as the LNB jitter. The reference oscillator signals are degraded by the additive noise of each amplifier in the signal path, as well as the fiber optic link.

3 Instrument Geometry

In order to meet the second requirement above, a four-element API array was chosen. The geometry of the four-element array is designed to

(a) Preserve the current 300 m nearly east-to-west (E-W) measurement,

(b) Add a 300m roughly north-to-south (N-S) measurement baseline, and

(c) Add a 1km measurement baseline that bisects the other two.

The location of the new antenna pads was constrained by the availability of land that is controlled by the VLA: the central square mile which houses the VLA facilities and staff, and by the density and location of existing buried infrastructure. The location of the two new antenna pads minimized conflict with existing buried lines, while providing the lengths and directions suitable for the new design.

The "east-to-west" (E-W) baseline is not strictly east-to-west; it is aligned with the VLA east arm, at an angle of approximately 20.5^{o} south of east. The "north-south" (N-S) baseline forms an angle of 86.3^{o} to this, so the near-orthogonality of these two baselines will allow the turbulence field to be decomposed into longitudinal and transverse components, regardless of the fact that the array is not cleanly aligned to compass points. In addition, wind patterns at the VLA site are such that the prevailing wind direction (taken to be the driving energy source for the

turbulent flow) is largely westerly about 70% of the time, and southerly about 30% of the time. The orientation of the measurement baselines is therefore such that the transverse-versus-longitudinal measurement is available at all times, and will allow for full recovery of the 2-dimensional turbulent field statistics.

Table 1 lists the physical and effective ("projected") baseline lengths and directions, and their E-W and N-S components. The inclusion of the projected length column accounts for the position of the satellite signal in the sky. If the satellite were at the zenith, the effective baseline would be equal to the projected baseline. But the elevation angle of the satellite foreshortens any projection in the north-south direction, and the azimuthal angle foreshortens the projection in the east-west direction. From the API geographical centroid, the satellite is located at an elevation angle of 49.75^{o} and an azimuthal angle of 168.34^{o}.

Table 3: API baseline vectors. "Length" is the distance measured along the ground. "Projected length" is that distance projected to the angles at the location of the satellite in the sky

Baseline	Length (m)	Projected length (m)	Bearing degrees from N	E-W component (m)	N-S component (m)
N - W	345	298	26.50	133	267
W – E	454	450	66.37	412	180
W – S	880	744	134.13	534	518
N – E	291	267	115.94	240	117
N – S	1038	813	152.63	374	722
E - S	823	629	164.81	165	607

If the turbulent field were azimuthally symmetric, the baseline direction would be arbitrary – all orientations would produce the same result. But if the theory of Tatarski holds, and distinct structure function coefficients exist for longitudinal and transverse fields, then the present API will tend to underestimate the turbulent effect much of the time, due to the intermediate angle it forms with the prevailing wind direction. The new system will provide measurement baselines parallel and perpendicular to the majority wind direction, and will span the space at all times, regardless of wind direction.

4 The Satellite: SES-1

Presently we are using SES-1, a television and communications satellite that orbits at 101^{o} W longitude in the geosynchronous Clarke belt. The advantage of using geosynchronous satellites as signal sources, as Radford [1996] mentions, is that they do not move across the sky (significantly), and therefore do not require receiving antennas with tracking capability. This greatly simplifies the complexity of the receiving stations.

SES-1 transmits two sinusoidal beacon signals, in addition to its payload of digital television signals. These beacon signals appear at 11.701 GHz and 12.901 GHz. Antenna technicians on

the ground use these beacon signals for aligning dishes toward the satellite. We use the 11.701 GHz beacon as the signal source for the API.

5 Elements of the Hardware Design

A description of the hardware that comprises the instrument subsystems will be followed by a description of the principles and techniques of their operation.

5.1 The Antennas

The antennas used to detect the satellite beacon signal are commercially-available dual-band (C/Ku) satellite communications dishes. These are model 1183 offset-feed paraboloid reflectors manufactured by Prodelin of North America, approximately 1.8 m in diameter, with a feed offset angle of 22.3°.

5.2 Front Ends: Receiver and 1st Frequency Conversion

The satellite beacon signal is extremely weak at the ground, typically picowatts. Commercially-available radio frequency small-signal amplifiers function optimally with signals in the microwatt (μW) to milliwatt (mW) regime. So the receiver electronics must provide 60 dB (a factor of one million) of gain to this signal while not significantly contributing self-generated noise. From the satellite's published Effective Isotropic Radiated Power (EIRP) specification and information about the transmitting antenna's beam pattern, we can calculate the link budget between the satellite-based transmitter and the ground-based receiver. The EIRP for a transmitting antenna is the product of the true transmitted power and the antenna gain factor.

The link can be characterized according to the Friis Transmission Equation, which gives the received power in terms of the transmitted power, transmitting antenna gain factor, receiving antenna gain factor, and path loss that is dependent on the distance R and wavelength λ:

$$P_R = \left(P_T + G_T\right) + G_R + 10\log_{10}\left(\frac{\lambda}{4\pi R}\right)^2$$

Expressed in decibels, EIRP is the sum of the transmitted power P_T and the transmitting antenna gain factor G_T. The published EIRP for SES-1 is 48 dBW, or approximately 63 kW. For λ = 2.6cm, R = 36,000 km, and G_R = 45 dB, this would yield P_R = -112 dBW = -82 dBm ~ 6.3 pW received at the ground.

The API front end is a receiver and low-noise block downconverter (LNB), Orbital Research LNB-5400X, which is a commercially-available device used in satellite communications and television groundstation applications. During the development phase of the API design, two competing products offered the right frequency coverage and the ability to accept an external frequency reference signal. Of these, the Orbital Research device offered the lower internal frequency jitter.

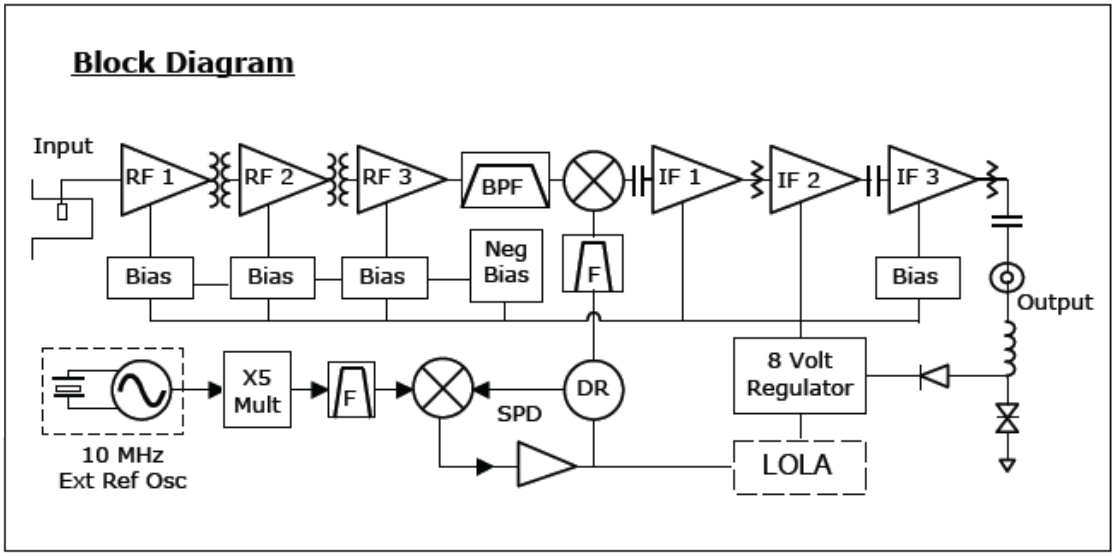

Figure 1: Block diagram of Orbital Research LNB-5400X, showing gain cascade, bandpass filtering, and frequency conversion. The significance of these will be discussed in detail in Section **Error! Reference source not found.** (Source: www.orbitalresearch.net)

The block diagram of this receiver is shown in Figure 1. The receiver receives and amplifies the 11.7 GHz satellite beacon signal. An internal voltage-controlled oscillator (VCO) phase locks to an external 10 MHz reference and produces a Local Oscillator (LO) signal at 10.75 GHz. The resulting 951 MHz Intermediate Frequency (IF) signal is amplified with a conversion gain of about 60 dB, which gives about -30 dBm, or ~1 μW at the output of the LNB. This signal power is within the dynamic range of relatively inexpensive, general-purpose commercially-available amplifiers to amplify without undue SNR degradation.

The front-end interface module houses a fiber optic receiver, which receives a 10MHz reference signal from the electronics shack. This signal requires about 40 dB of amplification, in two stages – a low-noise stage followed by a high-gain stage -- in order to drive the reference input of the LNB, which contains an integral oscillator that is phase-locked to the common 10 MHz signal. This VCO operates at 1.34375 GHz; the 8th harmonic of the VCO output is selected by a bandpass filter and used as the 1st-conversion LO signal, at 10.75 GHz.

The 951 MHz IF out of the LNB is bandpass filtered to 12MHz around the beacon frequency and amplified with an additional 30dB of gain. The output of this amplifier drives the modulator of the fiber optic transmitter, which is optimally modulated by a signal power of 1mW. This signal is transmitted to the shack for further processing.

5.3 Local Oscillator

A common 10 MHz reference signal is used to phase lock both the 1st conversion front-end LNB oscillator and the 2nd conversion synthesizer, and as the LO for the third frequency conversion. This reference signal is provided by a stable 10 MHz crystal oscillator in the LO module. This oscillator can be phase-locked to an external 10 MHz signal, but will self-center its frequency tuning in the absence of an external reference signal. Since this signal is common to all

antennas, all other oscillators in the system are mutually coherent, and any frequency instability in this central oscillator will cancel in phase comparisons.

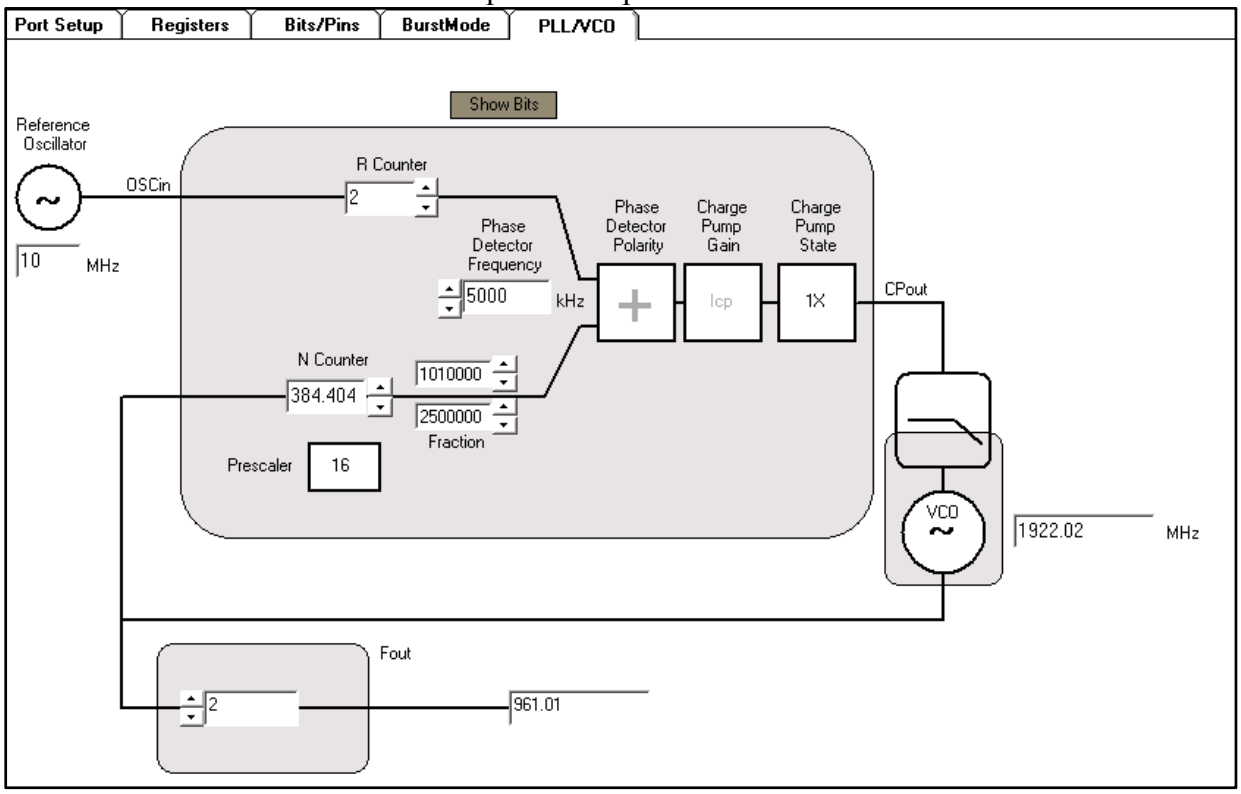

Figure 2: Screen capture of Texas Instruments LMX2531 CodeLoader software. This software allows graphical configuration of synthesizer and exports text-based hexadecimal commands that allow control of the device over a standard computer parallel port.

A digitally programmable VCO-based synthesizer (Texas Instruments/National Semiconductor LMX2531) produces the 2nd conversion LO signal at 961.01 MHz. This synthesizer is also phase-locked to the common 10 MHz reference signal. The synthesizer-programming interface allows independent control over output frequency, phase detector frequency, prescaler value, loop filter resistor and capacitor values, and charge pump gain. These parameters allow the performance of the synthesizer to be optimized for each selection of output frequency; for example, spurious output signals can be suppressed, but at the expense of increased noise floor. Figure 2 shows the graphical programming interface. This synthesizer is programmed for an output frequency 10 kHz away from the IF signal frequency. The synthesizer control register settings, as determined by the graphical controls, are stored to a text file where they can be read and asserted by the data acquisition computer.

The synthesizer parameters that influence the output frequency are

1. VCO frequency,
2. Phase detector frequency, and
3. Output frequency divider

These can be set independently, as long as the VCO frequency divided by the output frequency divider ratio is within the tunable range of the synthesizer (in the case of the LMX25312080, this

is 1904MHz < f_{VCO} < 2274MHz.) Additionally, the phase detector frequency, charge pump gain, and loop filter compoenent values can all influence spectral shape, the presence of spurious signal components, and overall oscillator jitter. From the LMX2531 datasheet, for phase detector frequency f_{PD} and oscillator (VCO) frequency f_{OSCin}, $f_{PD} = f_{OSCin}/R$:

> *Choosing R = 1 yields the highest possible phase detector frequency and is optimum for phase noise, although there are restrictions on the maximum phase detector frequency which could force the R value to be larger. The far out PLL noise improves 3 dB for every doubling of the phase detector frequency, but at lower offsets, this effect is much less due to the PLL 1/f noise. Aside from getting the best PLL phase noise, higher phase detector frequencies also make it easier to filter the noise that the delta-sigma modulator produces, which peaks at an offset frequency of $f_{PD}/2$ from the carrier.*

Similarly, the charge pump gain affects both the overall phase noise and the level of spurious signal components:

> *Increasing the charge pump current improves the phase noise about 3 dB per doubling of the charge pump current, although there are small diminishing returns as the charge pump current increases. From a loop filter design and PLL phase noise perspective, one might think to always design with the highest possible phase detector frequency and charge pump current. However, if one considers the worst case fractional spurs that occur at an output frequency equal to 1 channel spacing away from a multiple of the f_{OSCin}, then this gives reason to reconsider. If the phase detector frequency or charge pump currents are too high, then these spurs could be degraded, and the loop filter may not be able to filter these spurs as well as theoretically predicted. For optimal spur performance, a phase detector frequency around 2.5 MHz and a charge pump current of 1X are recommended.*

The five API 2nd LO frequencies that have been used since the beginning of the project – 961.005, 961.01, 961.015, 961.016, and 961.02MHz -- have all used f_{PD}=5MHz and charge pump gain = 1.

Once a new output frequency has been chosen and the relevant synthesizer parameters determined, the LMX2531 control register values can be exported from the CodeLoader program into a text file (include the output frequency in the name of the text file, e.g. synthregisters_961_016.txt). These register values are stored in the columns of a table which are indexed by the frequency select command, issued from the M360 synth_freq_cmd control point.

An example spectrum of a real oscillator signal is shown in Figure 3. These are the measured spectra of the API 2nd LO signal. In each successive frame of Figure 3, the frequency span is reduced by a factor of 10. At a sufficiently large span, 50MHz, the signal appears impulse-like (a), with some measurement noise. In (b), at 5MHz width, there is some apparent widening of the signal, though no meaningful spectral structure is resolved. In (c) at 500kHz width, other spectral features begin to emerge. These details were lost at the frequency resolution of the previous frames. Finally, in (d), 50kHz span, we see distinct features – other impulses, plus some structure in the noise floor. The 50kHz span is significant, because this is the sampling frequency of the API digitizer – any signal energy contained in this bandwidth will be sampled and will contribute to the total sampled signal energy.

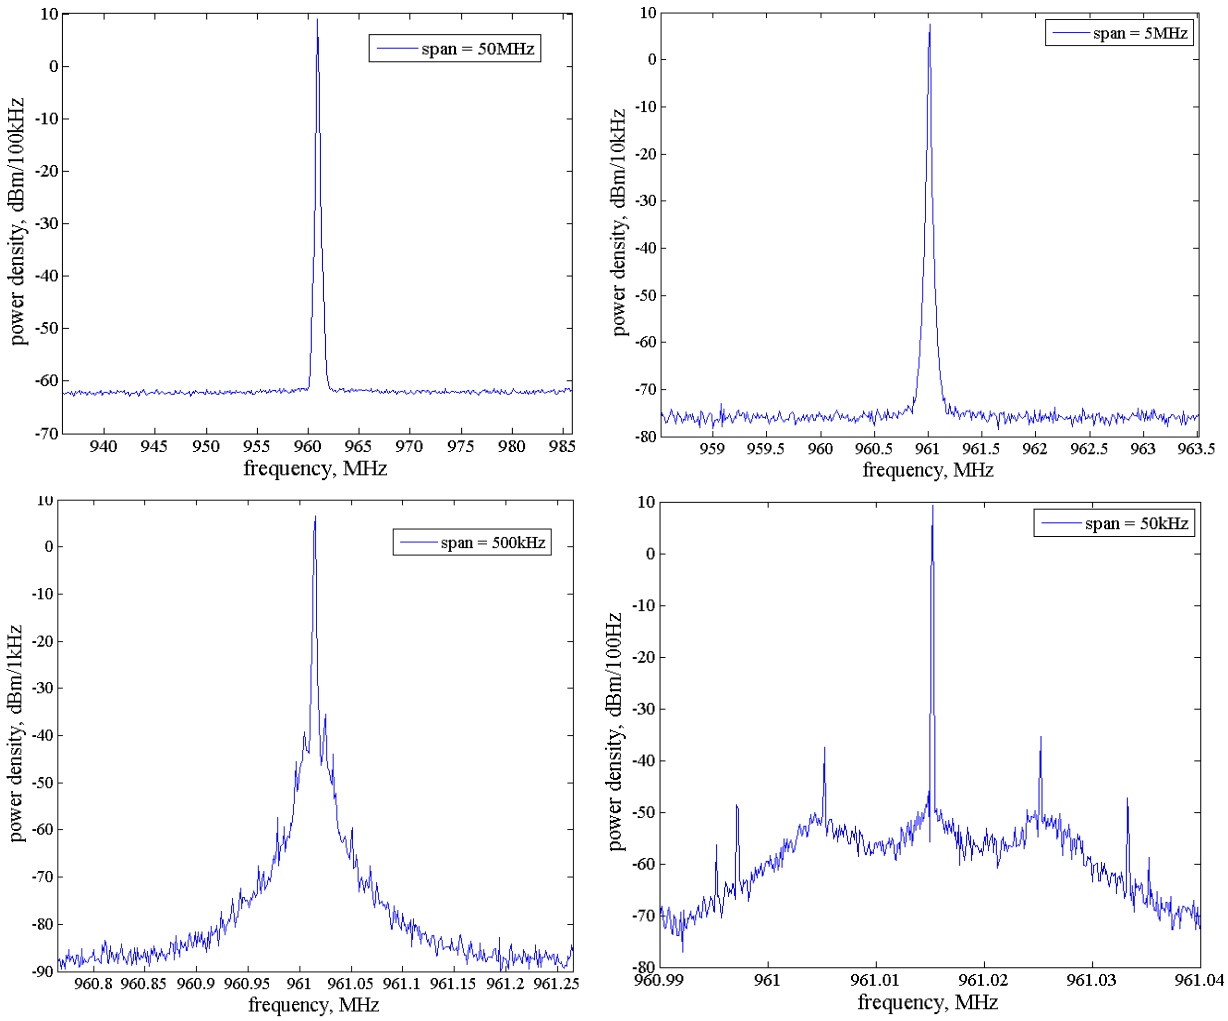

Figure 3: The LO signal at 961MHz, viewed at four different frequency scales. Each is measured using a constant number of frequency bins.

5.4 Fiber Optic Distribution of Signals

Fiber optic cable is used to distribute the 10 MHz LO reference signal from the shack to the antennas, and the 951 MHz IF signals from the antennas to the shack. A separate set of fiber optic cables carries Ethernet monitor and control signals between the shack and the VLA control building. These are 12-conductor bundles of double-armored single-mode fiber at an operating wavelength of 1310nm, buried to a depth of approximately four feet. This burial depth was found to provide optimal temperature stability at a reasonable labor cost.

The transmitter and receiver for the Local Oscillator reference signal are the TimeLink DL-series from Linear Photonics, optimized for distributing time and frequency reference signals over long distances. The transmitter/receiver pairs for the analog RF signals are Fiberspan AC123-T/R, which are standard parts commonly used in the distribution of cable television signals. The transmitter is a direct-modulation 1310 nm laser driven by the incoming 951 MHz signal at a power level of approximately 0 dBm. The receivers are photodiodes followed by an RF

amplifier cascade with a bandwidth of approximately 1.5 GHz. The receivers impart noise to the analog signals, particularly at low (< 1 MHz) frequencies. It is therefore necessary to bandpass filter the RF output of each fiber link.

The 1 km haul of the longest API baseline made the coaxial copper distribution of signals impractical, due to signal attenuation in the cable, which is a function of cable length. Using premium low-loss cable, the loss over a 1 km distance at 951 MHz was approximately 72 dB. This was deemed unacceptable, requiring fiber optic cable to be used for the longest baseline. In the interest of standardizing the four antennas' signal performance, regardless of distance, fiber optic links are used to carry the 10 MHz local oscillator reference to each antenna, and to carry the 1st IF signal from each antenna to the processing shack. This turned out to be one of the most challenging aspects of the project, owing to the jitter contributed by the laser modulator.

The fiber optic link operates by modulating an RF signal onto a directly-modulated laser, guiding the laser signal through a fiber optic cable, and demodulating or detecting the signal with a photodiode and recovering the RF signal. For the API distances, the signal attenuation in the fiber optic link is dominated by fiber splices and connectors, and by the modulator/demodulator circuits, and the loss is approximately the same for all four antennas, independent of distance.

The API uses two distinct models of fiber optic link: a low-jitter (and relatively expensive) transmitter/receiver pair for the 10 MHz LO reference, Linear Photonics TimeLinkDL DTMT(R)HFMC, and a lower-cost pair, Fiberspan AC-123T/R for the 1st IF signal. The IF signal fiber links were developed and marketed for the distribution of cable television signals, but were adapted for use in the VLA (and ultimately the API) signal distribution system.

Originally all the fiber components were the Fiberspan models, which are used successfully in the VLA for local oscillator distribution. In fact, the units used in the API are surplus stock from the VLA upgrade project. But these were found to add an unacceptable amount of jitter to the 10 MHz signal. The VLA reference signals are high enough in frequency that they do not get corrupted by the frequency response of the modulator/demodulator link. The TimeLink models are specifically designed to transmit and receive timing reference signals, such as the API 10 MHz signal, and do not add appreciable jitter to the signals.

Table 4 shows the RMS jitter, from 1 to 100Hz, of the 10MHz oscillator alone, and after being detected on each of the TimeLink and Fiberspan links. Also listed is the RMS jitter of the LMX2531 synthesizer at 961MHz.

Table 4: Jitter of the 10MHz reference, in isolation and with each of two fiber optic distribution systems.

Source	**Jitter (ps)**
10MHz Wenzel oscillator (reference and 3rd LO)	0.402
Fiber link (Timelink)	0.405
Fiber link (Fiberspan)	4.965

5.5 2nd Frequency Conversion

After the IF is demodulated from the fiber, it is amplified by about 40 dB before entering the 2nd conversion frequency mixer. The 961.01 MHz LO signal converts the IF to a nominal 10.01 kHz

signal. The 2nd IF signal is bandpass filtered by a narrow (+/- 500 kHz) filter to eliminate spurious signals and fiber optic link artifacts that would otherwise corrupt the signal.

5.6 3rd Frequency Conversion and Baseband Conditioning

The 10.01 MHz 2nd IF is mixed with a copy of the central 10 MHz LO frequency reference, resulting in the desired nominal baseband signal of 10 kHz. The baseband signal is low-pass filtered by a programmable active filter circuit with a cutoff frequency of 20 kHz, to prevent aliasing in the digitizer, which has a Nyquist frequency of 25 kHz. The baseband filter can be reconfigured with different cutoff frequencies and gains, or as an entirely different filter topology, should the signal require different filtering.

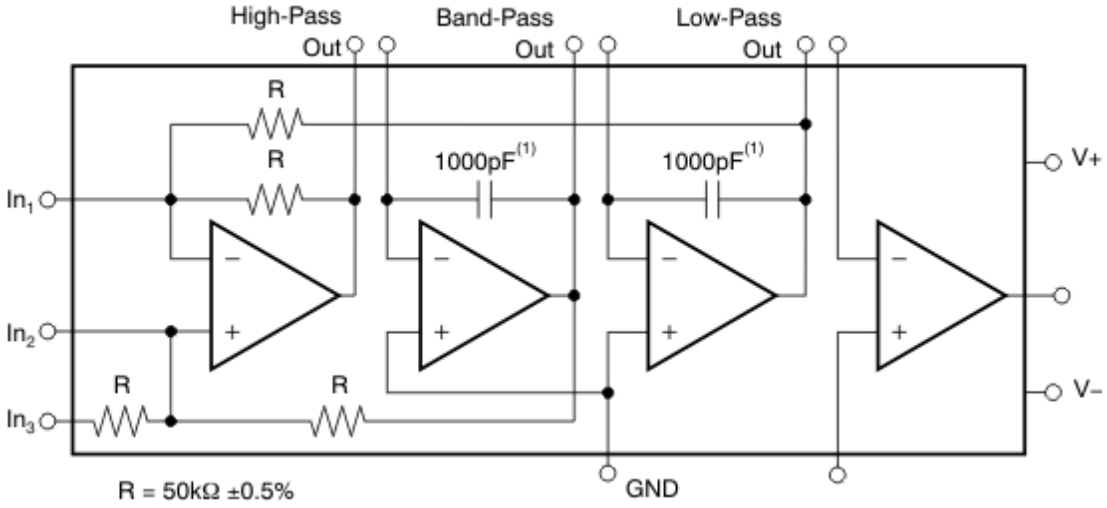

Figure 4: Schematic representation of UAF42 programmable baseband filter.

The baseband lowpass filter is realized with a Texas Instruments UAF42 biquad filter, whose schematic representation is shown in Figure XX. This filter can be configured using FILTER42 program from Texas Instruments. This program allows a filter topology and set of characteristics to be specified in a GUI and returns the component values and pin connections necessary for realization of that filter.

5.7 Digitizer

The digitizer resides in the Data Acquisition (DAQ) module, and is an APEX Systems model STX-104 PC/104-standard multiplexed-input, 16-bit analog-to-digital converter (ADC) that can accommodate up to 16 input channels, for a maximum sample rate of 200 ksps. In this configuration, we use four input channels, each with a sample rate of 50 ksps. The inputs are configured as differential pairs, although each signal is single-ended. This pseudo-differential arrangement has the advantage of cancelling noise that is common to each signal cable and its cable shield, and allows a broader range of adjustable gain within the digitizer module. Each channel can sample a +/-10V amplitude signal, although for smaller-level signal, a reduced range is preferable, so that the signal is represented by as many bits as possible. Each channel is configured with the maximum possible gain, which allows +/- 1.25V full scale amplitude.

The 200 kHz digitizer sample clock is the one reference signal in the system that is not coherent to the 10 MHz reference oscillator. This is acceptable, since any drift in this clock will appear in common to all four channels, and therefore it does not impart any drift onto the sampled signals.

6 Data collection and analysis

The API instrument produces raw voltage streams from the two antennas, which are then collected, correlated, integrated, and processed into the phase difference time series (one phase difference per second) in software running on a Module Interface Board (MIB) unit located in the electronics shack. The phase difference time series is then analyzed in 10-minute segments in a program deployed on one of the VLA monitor and control remote server computers. This second program derives the interesting structure function quantities (power-law slope, corner time, and saturation phase) and writes them to the VLA archive database. This software all pre-dated this thesis project, having been developed by VLA staff for the original 1997 API installation, and was taken as-is from the existing implementation for the new API.

6.1 Integration and Correlation

The MIB CPU requests a data packet from the DAQ module once per second. This packet contains nominally 200,000 samples, 50,000 each from the four channels. Although we have not yet installed the third and fourth antennas or their associated electronics modules, the data acquisition system is capable of processing all four potential channels. The program forms the six pairs of signals, places each pair in quadrature, and correlates.

The complex correlation of two time series $x_1(t)$ and $x_2(t)$ is given by

$$y(t) = \int_0^t x_1(t+\tau)x_2^*(\tau)d\tau$$

where the conjugation is accomplished by Hilbert transforming x_2,

$$x_2^*(t) = x_2(t) \otimes \frac{1}{i\pi t}$$

In the API processor, the integral becomes a summation over discrete values of time:

$$y(k) = \sum_{n=1}^{50,000} x_1(n+k)x_2^*(n)$$

The one-second data from each pair is correlated in the data acquisition module, resulting in the real (Re) and imaginary (Im) parts of the correlation. The phase is calculated as

$$\phi = \arctan\left(\frac{\mathrm{Im}(y)}{\mathrm{Re}(y)}\right)$$

and communicated over the network to the monitor and control server.

6.2 Derived Structure Function Quantities

1. The remote server program operates on a 10-minute stream of one-second phase difference data. It performs the following operations on the time series:

2. Unwrap the phase to remove 2π phase ambiguities
3. Calculate and remove a least-squares quadratic fit, to remove phase drift due to satellite motion
4. Calculate the saturation phase, which is the RMS of the residual from (2)
5. Calculate the structure function of the phase difference time series

For more details, see Butler and Desai [1999].

6.3 Baseband signal frequency calculation

One further calculation is made in software on the MIB, which is of use to the API instrument. Once per ten minutes, the MIB software performs a Fast Fourier Transform (FFT) on one second's worth of sampled signal from one of the four channels. The program removes the DC component and identifies the frequency at which the maximum occurs in the spectrum and displays and archives this value. Because the master oscillator in the API system is not synchronized to the oscillator on board the satellite, the two oscillators are free to drift in frequency with respect to each other, and the result of this mutual drift is that the baseband signal may drift out of range of its lowpass filter, and ultimately of the digitizer. Periodic adjustments are made to the 2^{nd} LO signal to compensate for this drift.

6.4 Data Visualization

Presently, the API Engineering Screen (Figure 5) displays voltages, power levels, temperatures, and other diagnostic data in table form, updated continuously by the M&C system. In addition, the Weather Screen (Figure 6) displays the graphical time series of the ten-minute RMS phases from each of the six baseline products. A proposed future GUI screen would show the RMS phase value for each baseline superimposed on a map of the API array, as in Figure 7.

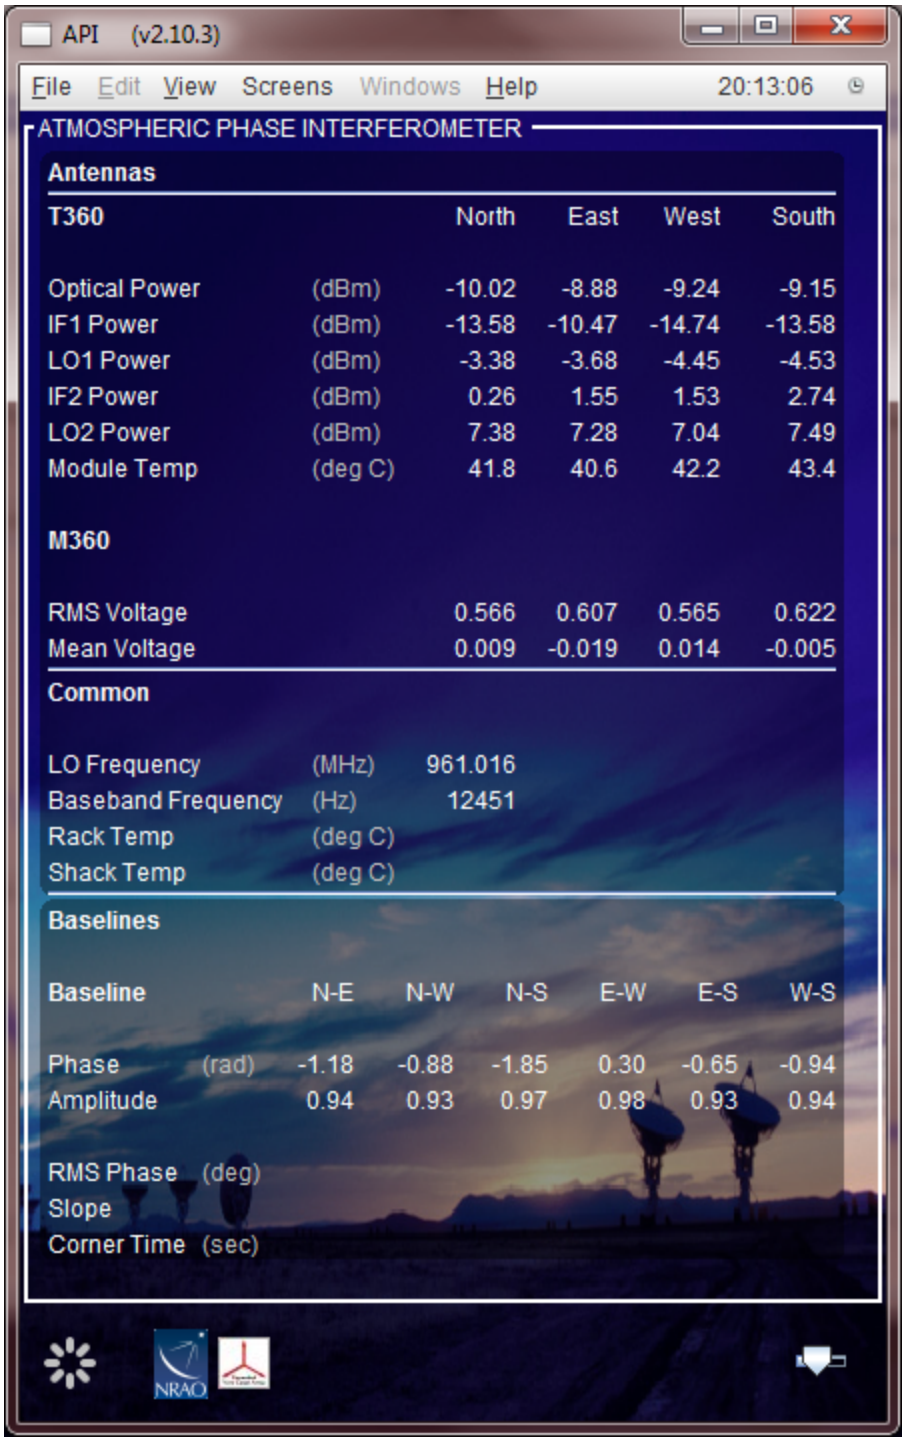

Figure 5: API Engineering screen, showing diagnostic monitor points for the electronic modules throughout the system.

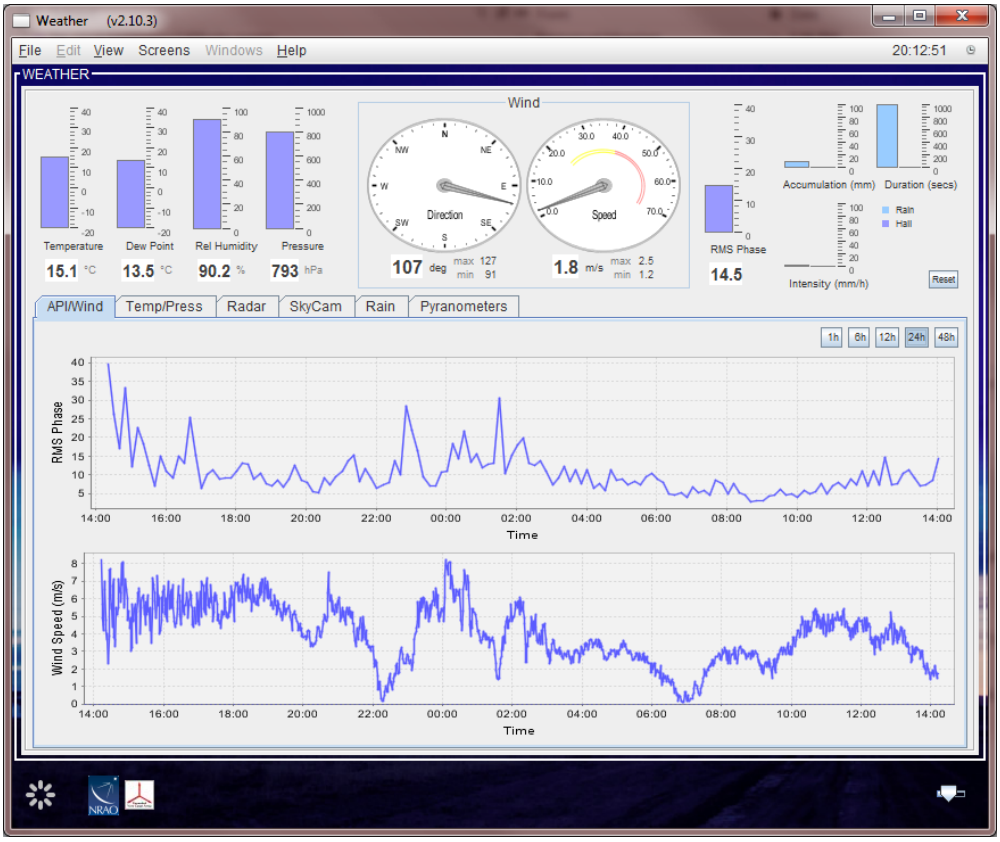

Figure 6: Weather screen API/Wind tab, showing 24 hours of API RMS phase and wind speed.

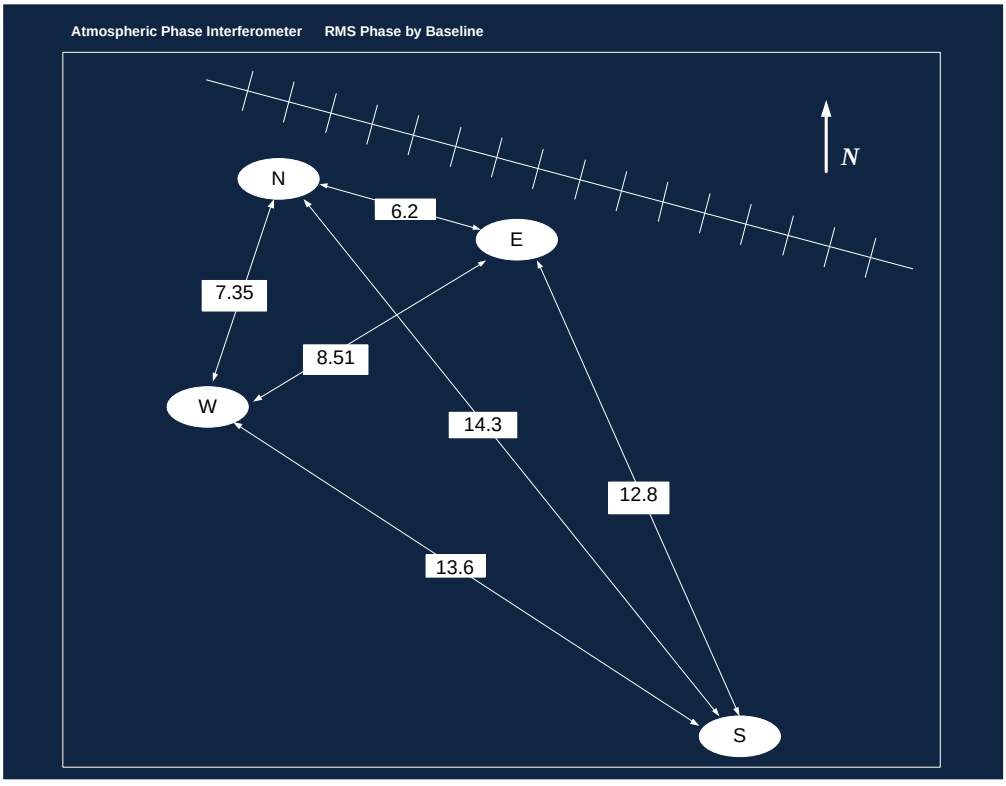

Figure 7: API RMS phase map (proposed future GUI display)

7 An Example Analysis of API Data

It will be illustrative to look at an example of the API measurement products at this point. Figure 8 shows the raw phase difference from the existing API, sampled once per second, spanning a period of about one day. It is apparent that the phase wraps many times around the phase circle over this period due to 2π ambiguity, motivating the phase-unwrapping step of the analysis software.

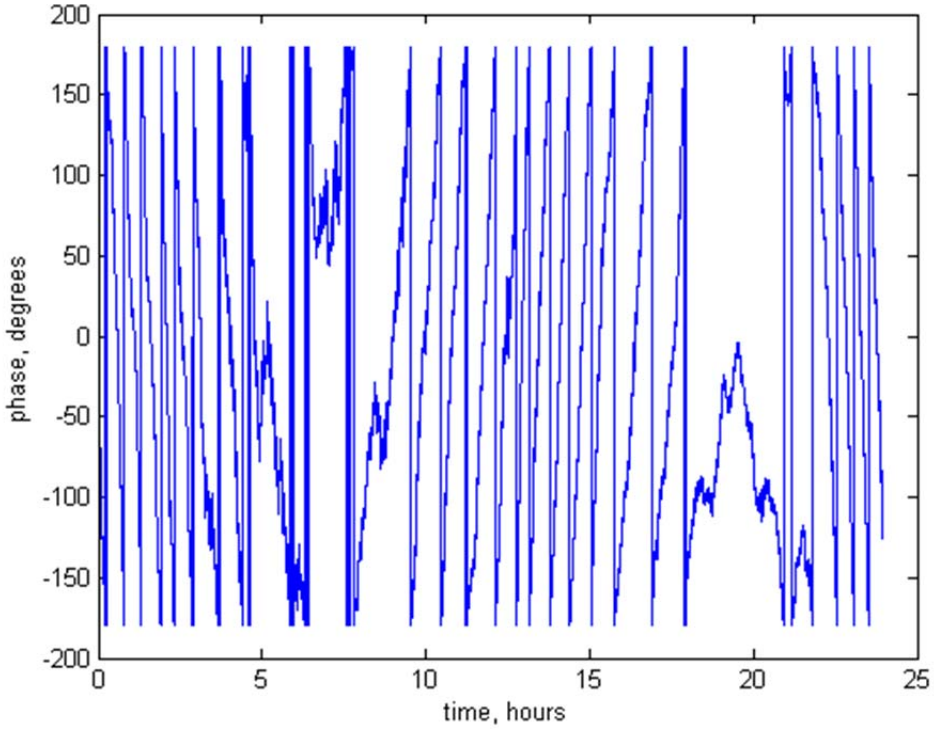

Figure 8: Raw phase time series for one day, measured once per second. Several phase wraps occur due to the Doppler frequency variation of the moving satellite.

Figure 9 shows the same raw phase data, after unwrapping. Even after the unwrapping, there is almost 4000 degrees of phase excursion, which is almost entirely due to the effect of the motion of the satellite within the beam of the antennas. Figure 10 shows the ephemeris of the satellite throughout one day. The visible discontinuity in the position course requires that the satellite perform a daily position correction, and this discontinuity is evident in the frequency of the detected and converted "baseband" signal, as shown in Figure 11.

This large change in frequency is the result of the frequency of the beacon signal changing according to the Doppler formula,

$$f_{observed} = \left(\frac{c + v_r}{c + v_s} \right) f_0$$

where c is the speed of propagation of the signal through the medium (taken to be the speed of light in air), v_r is the velocity of the receiver, v_s is the velocity of the source, and f_0 is the nominal

signal frequency. The observed frequency will be greater than the nominal frequency when the source is approaching and less than the nominal frequency when the source is receding.

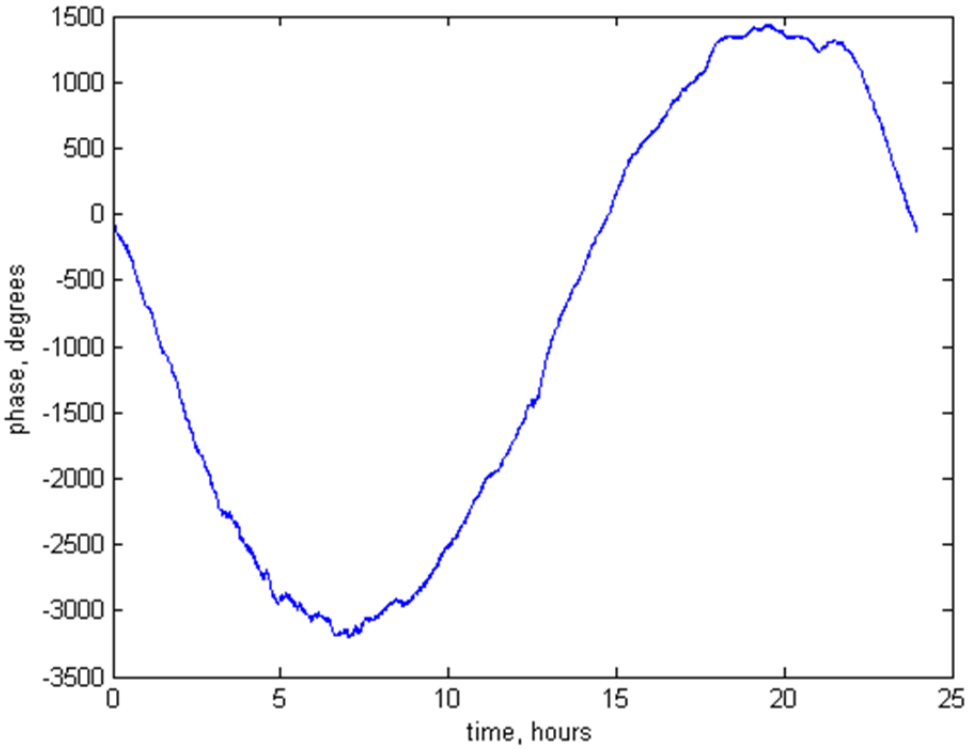

Figure 9: Same raw phase data, unwrapped to show complete phase excursion.

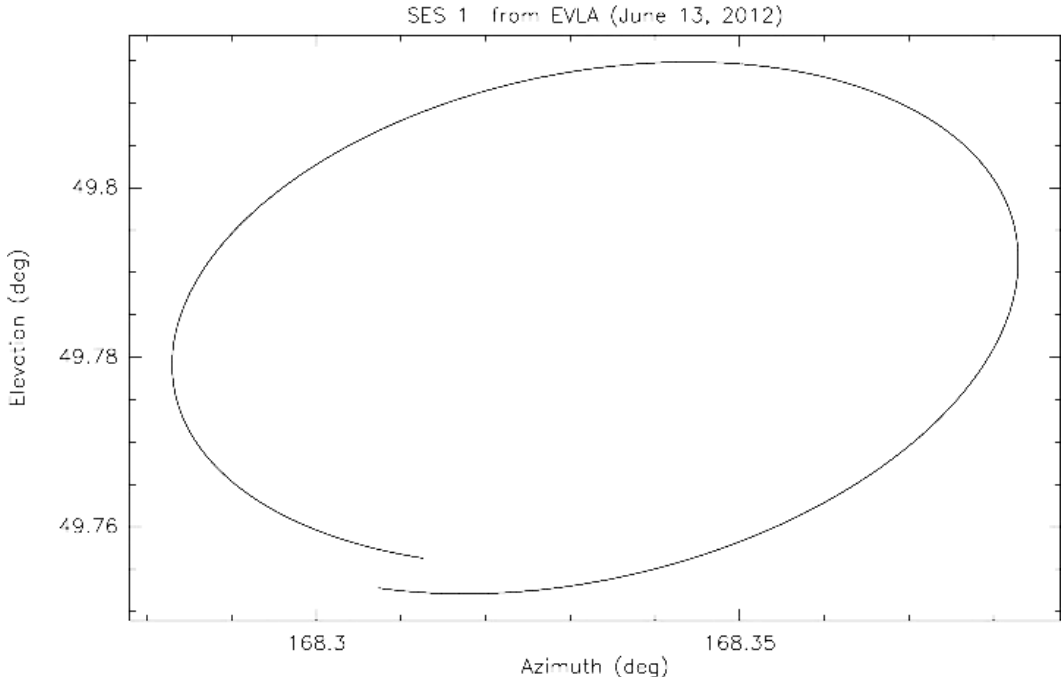

Figure 10: Satellite SES-1 ephemeris for one day

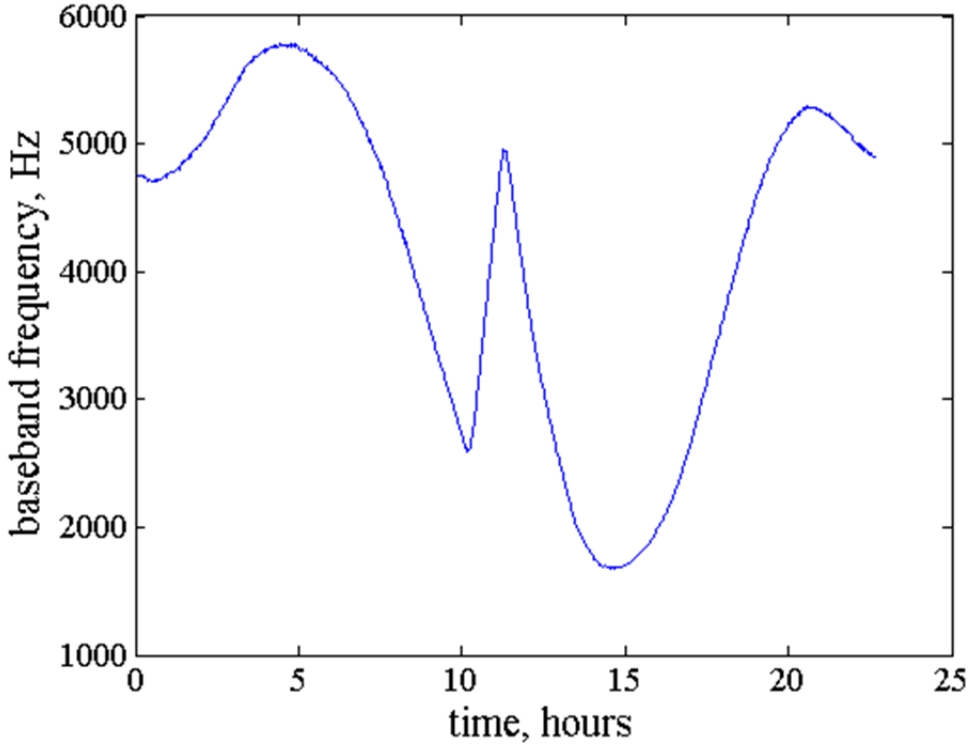

Figure 11: Signal baseband frequency, showing Doppler variation from satellite motion and discontinuity from position correction.

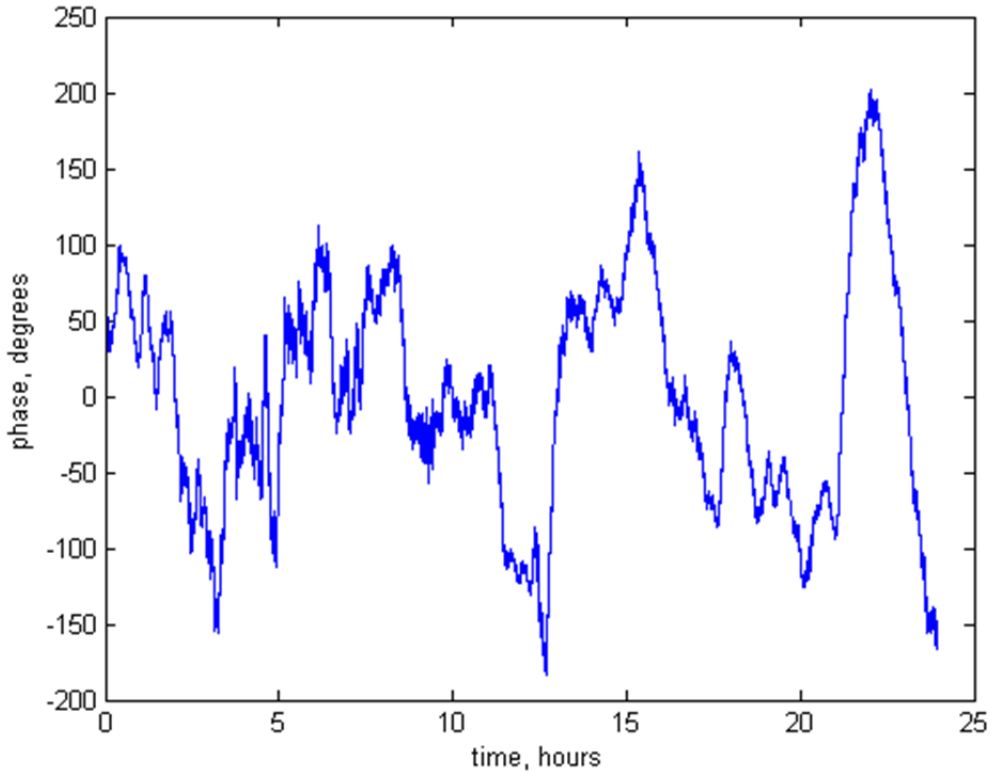

Figure 12: One day of raw phase data, diurnal sinusoidal trend removed

It is difficult to discern at the scale of Figure 9, but the thickness of the line varies throughout the day. This is due to the changing atmospheric stability throughout the day, and it is this quantity that we seek to measure. Figure 12 shows the same data, but with the diurnal sinusoid removed; the variance on timescales of minutes and hours becomes more apparent. Notice that there are ten-minute regions of Figure 12 that would not be well-fit by a quadratic curve, for example regions where there are multiple inflection points. This will naturally influence the RMS measurement and except in rare cases that condition is unlikely to persist for multiple contiguous ten-minute periods.

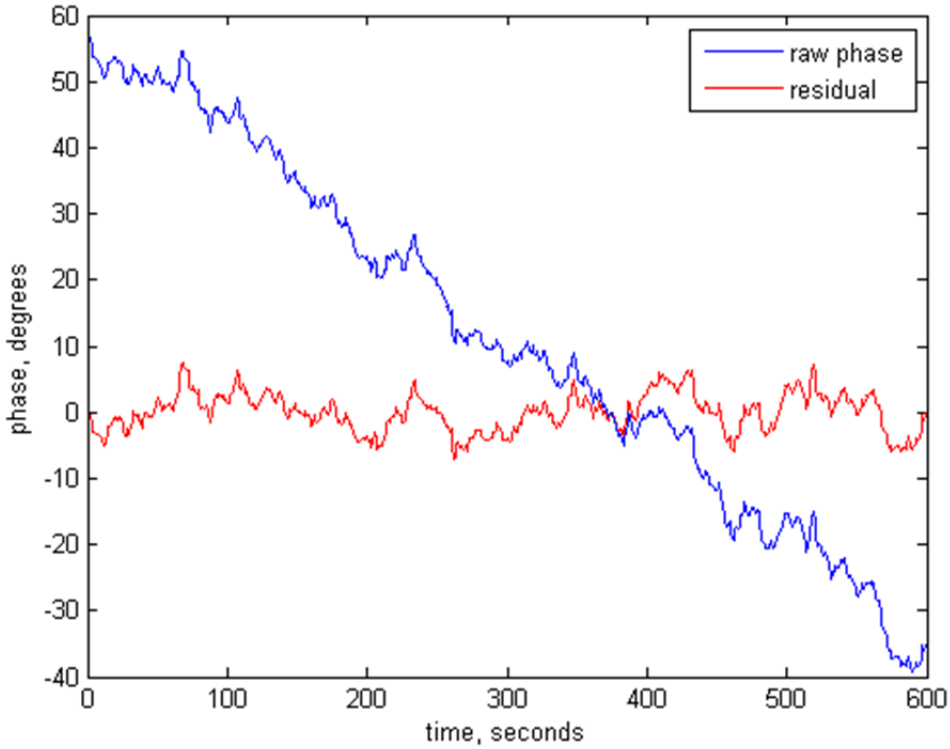

Figure 13: Ten minutes of raw phase (blue) and residual (red) after quadratic fit. The RMS of the red trace is 2.94 degrees.

Figure 13 shows the raw phase values for the final ten-minutes of the data in blue, and the same data, but with a quadratic fit removed, in red. The quadratic trend accounts for the satellite motion and slow temperature-induced phase drift.

Calculating the root-mean-square of the residual from that fit gives us the API measurable – the saturation phase, in this example 2.94 degrees. We can perform this exercise for each of the 10-minute segments of the entire data set for the day, deriving a saturation phase for each of these segments. Figure 14 shows this quantity for the entire 24-hour data set. As noted above, the software also determines the power-law exponent of the structure function for each of the 10-minute segments – Figure 15 shows that quantity for each ten-minute period plotted for the entire 24-hour data set.

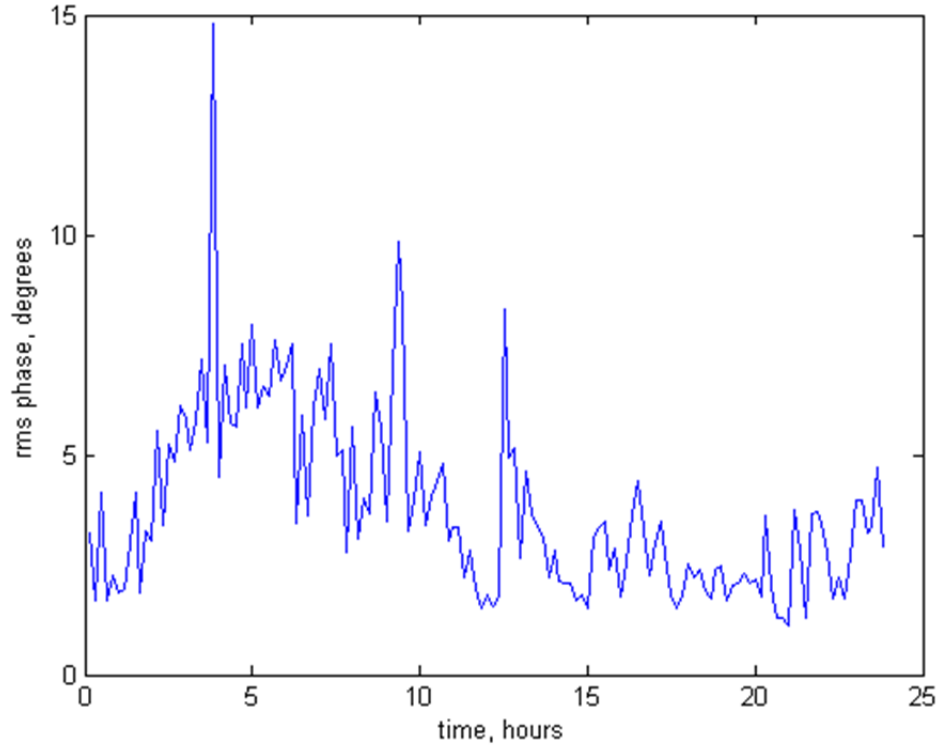

Figure 14: Time series of saturation phase derived each 10 minutes over the entire 24 hours of data shown in Figure 8.

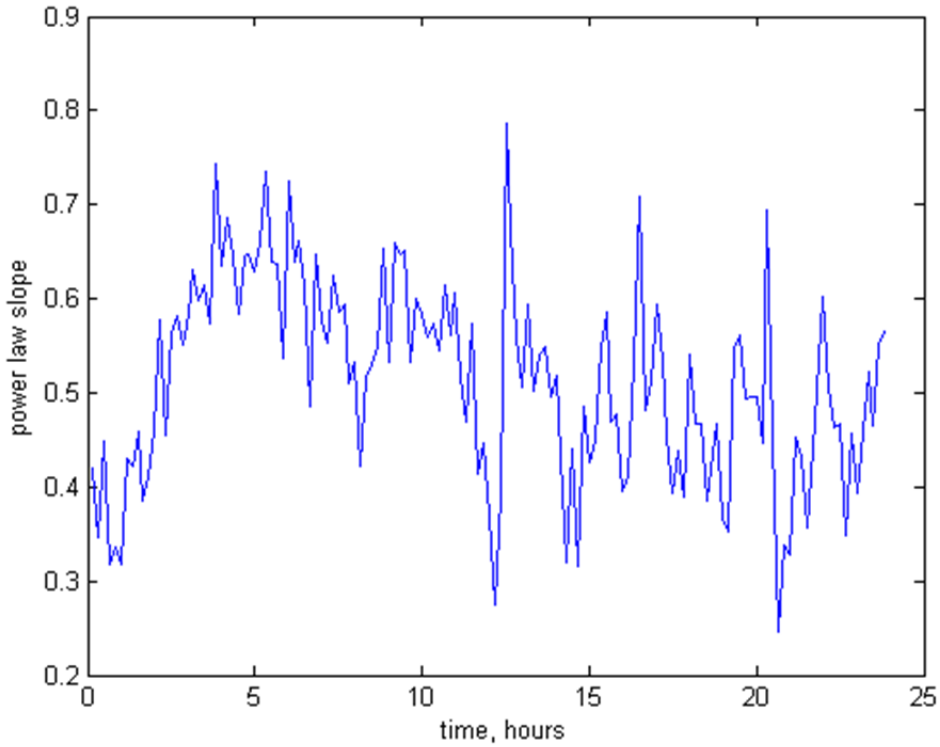

Figure 15: Time series of power-law exponent (slope) derived each 10 minutes over the entire 24 hours of data shown in Figure 8.

The VLA Monitor and Control (M&C) system provides both real-time monitor data – for the API, these data include signal power levels, circuit board temperatures and supply voltages, and more importantly the instantaneous one-second phase and correlated amplitude. In addition, these quantities are archived in a database where retrieval of historic data allows matching of trends, for example outdoor air temperature, windspeed, humidity, or insolation, to signal levels and phase. This archived database has enabled data analysis that otherwise would have required costly dedicated data-logging systems to be deployed in advance. Much of the progress that has occurred in the API project has resulted from discovering unanticipated interactions that could only have been revealed by inspection of this comprehensive data archive.

8 Future Work

Several further improvements to the Phase I design are planned, which are expected to improve overall phase stability of the instrument, and result in lower instrumental noise and variation, and therefore better overall instrumental sensitivity. In particular, work is underway to improve the jitter performance of the Local Oscillator distribution, and to thermally stabilize the electronics shack and its components.

Of all the components comprising the API, the Fiber optic and Teflon-dielectric coaxial cables (because of their low thermal mass) are particularly sensitive to rapid temperature changes. Thus, temperature changes on the scale of one to six hundred seconds result in an electrical or optical length variation, which translates to induced phase variation, of a sufficient magnitude to affect the API measurement.

The HVAC unit in the electronics shack occasionally cycles between its set point and its hysteresis point, when its cooling capacity is out of balance with the heat load of the shack. This cycling provides the rapid temperature variation to which the API cables are sensitive. To mitigate this effect, we are reinforcing the insulation of the shack, thermally isolating sensitive components from air that is changing temperature quickly, and balancing the heat load so that the HVAC runs continuously. These three parallel approaches will minimize the instrumental sensitivity to rapid temperature variation.

More comprehensive data visualization tools may be required as NRAO staff become more acquainted with the new instrument's data output. Also, further analysis of API and weather station data can probe whether the direction-dependence of the structure function constant as predicted by Tatarski is apparent.

9 Acknowledgements

Much of the substance of this memo was taken from my Master's thesis, which available through the NRAO Library and the NM Tech Library. That work represented the combined effort of many people, including my advising committee: Bryan Butler, Dave Westpfahl, Richard Sonnenfeld, and Sharon Sessions, as well as the contributions from all who worked and collaborated on the project: Eric Chavez, Steve Tenorio, Brent Willoughby, Hichem Ben Frej, as well as all who have labored on this problem before me.